arXiv:1603.04139v2 [cs.CV] 28 Sep 2016

SSSC-AM: A Unified Framework for Video Co-Segmentation by Structured Sparse Subspace Clustering with Appearance and Motion Features

Junlin Yao* Frank Nielsen†

September 30, 2016

Abstract

Video co-segmentation refers to the task of jointly segmenting common objects appearing in a given group of videos. In practice, high-dimensional data such as videos can be conceptually thought as being drawn from a union of subspaces corresponding to categories rather than from a smooth manifold. Therefore, segmenting data into respective subspaces — subspace clustering — finds widespread applications in computer vision, including co-segmentation. State-of-the-art methods via subspace clustering seek to solve the problem in two steps: First, an affinity matrix is built from data, with appearance features or motion patterns. Second, the data are segmented by applying spectral clustering to the affinity matrix. However, this process is insufficient to obtain an optimal solution since it does not take into account the *interdependence* of the affinity matrix with the segmentation. In this work, we present a novel unified video co-segmentation framework inspired by the recent Structured Sparse Subspace Clustering (S^3C) based on the *self-expressiveness* model. Our method yields more consistent segmentation results. In order to improve the detectability of motion features with missing trajectories due to occlusion or tracked points moving out of frames, we add an extra-dimensional signature to the motion trajectories. Moreover, we reformulate the S^3C algorithm by adding the affine subspace constraint in order to make it more suitable to segment rigid motions lying in affine subspaces of dimension at most 3. Our experiments on MOViCS dataset show that our framework achieves the highest overall performance among baseline algorithms and demonstrate its robustness to heavy noise.

Keywords: Object video co-segmentation, subspace clustering, self-expressiveness model, sparse representation, affine motion subspaces.

1 Introduction, prior work and contributions

Over the past several decades, we have witnessed an explosion in the availability of video data. Just to give a figure, let us say that the amount of new videos uploaded to YouTube every minute is estimated to be 300 hours! Segmenting videos into multiple spatio-temporal areas, and extracting useful information effectively and computationally efficiently from them becomes a crucial issue and has broad applications, such as 3D reconstruction, action recognition, etc. Various methods have been proposed to deal with this challenging video segmentation problem. For example, subspace clustering based methods [20] seek to discover low-dimensional representation of high-dimensional data points.

Co-segmentation [8, 23, 9] consists in segmenting simultaneously multiple images or videos into several subregions denoting common objects. Thus co-segmentation searches for objects jointly appearing in videos, and provides additional information in the absence of supervisory information [9].

*École Polytechnique, Palaiseau, France.
†École Polytechnique, Palaiseau, France. Corresponding author: Frank Nielsen (email: Frank.Nielsen@acm.org)

1.1 Previous work

1.1.1 Subspace clustering

Many subspace clustering algorithms have already been proposed, including algebraic algorithms [19], iterative methods [1], statistical methods [7], and spectral clustering based methods [21, 10, 17]. Among these algorithms, the Structured Sparse Subspace Clustering (S^3C for short) algorithm [11] outperforms all others due to its framework promoting coherence between the affinity matrix and segmentation, as well as its robustness with respect to noise and outliers. However, S^3C only considers linear subspaces while it is frequent to encounter data lying in a union of *affine* rather than linear subspaces for dynamic data-sets. For instance, the motion segmentation problem consists in clustering data drawn from 3-dimensional affine subspaces. Though there exists a naive way that data coming from affine subspaces can be treated as if they lie in linear subspaces, this method may be potentially unable to distinguish subspaces from each other [6]. Since video co-segmentation involves motion segmentation and thus affine subspaces, it is necessary for the subspace clustering method to have the ability to deal with data points lying in affine structures.

1.1.2 Video co-segmentation

Several methods, such as [5, 22], have been proposed to tackle the video co-segmentation problem. Both of the methods, similar to our framework, seek to jointly segment objects in video sequences in the absence of supervisory information. In [5], Chiu et al. formulate a generative multi-video model to enable multi-class video co-segmentation. In the presence of noisy motion information, such as objects moving together or existence of similar motions, its segmentation performance degrades quickly. In [22], Wang et al. propose a co-segmentation algorithm given facts exploited by subspace clustering based method called *Low-Rank Representation* (LRR). Both their method and ours conceptually view video data as being drawn from a union of subspaces. Compared to their assumption of background/foreground binary classification, our framework assumes a multi-class structure of videos and yields multi-class segmentation results as in [5]. Also, their algorithm suffers from limitations that it is unable to tackle defocus blur, noise and outliers. Inspired by those recent works, our framework seeks to better cope with videos corrupted by noise and outliers, as demonstrated by our experiments, where [5, 22] serve as a benchmark.

1.2 Paper contributions

In this paper, we propose a *unified* framework for video co-segmentation that enables to obtain *better coherent segmentation* results. We summarize our contributions as follows: First, we improve the detectability of motion trajectories corrupted by missing data using extra-dimensional signatures. Second, in order to deal with motion features lying in 3-dimensional affine subspaces, we reformulate the structured sparse subspace clustering by adding the affine subspace constraint, and solve this novel optimization problem efficiently using the Alternating Direction Method of Multipliers (ADMM) method. By integrating this subspace clustering approach, we propose a unified optimization framework, dubbed Structured Sparse Subspace Clustering with Appearance and Motion Features ($S^3C - AM$).

1.3 Paper organization

In Section 2, we introduce the subspace clustering problem and the S^3C algorithm for clustering data points lying in a union of linear subspaces. In Section 3, we motivate and formulate the unified framework for video co-segmentation. In Section 4, we study experimentally the effectiveness our framework on real video data sets. In Section 5, we discuss on some possible extensions of our work. Finally, Section 5 concludes the paper. The notations used throughout the paper are concisely recalled in Appendix A.

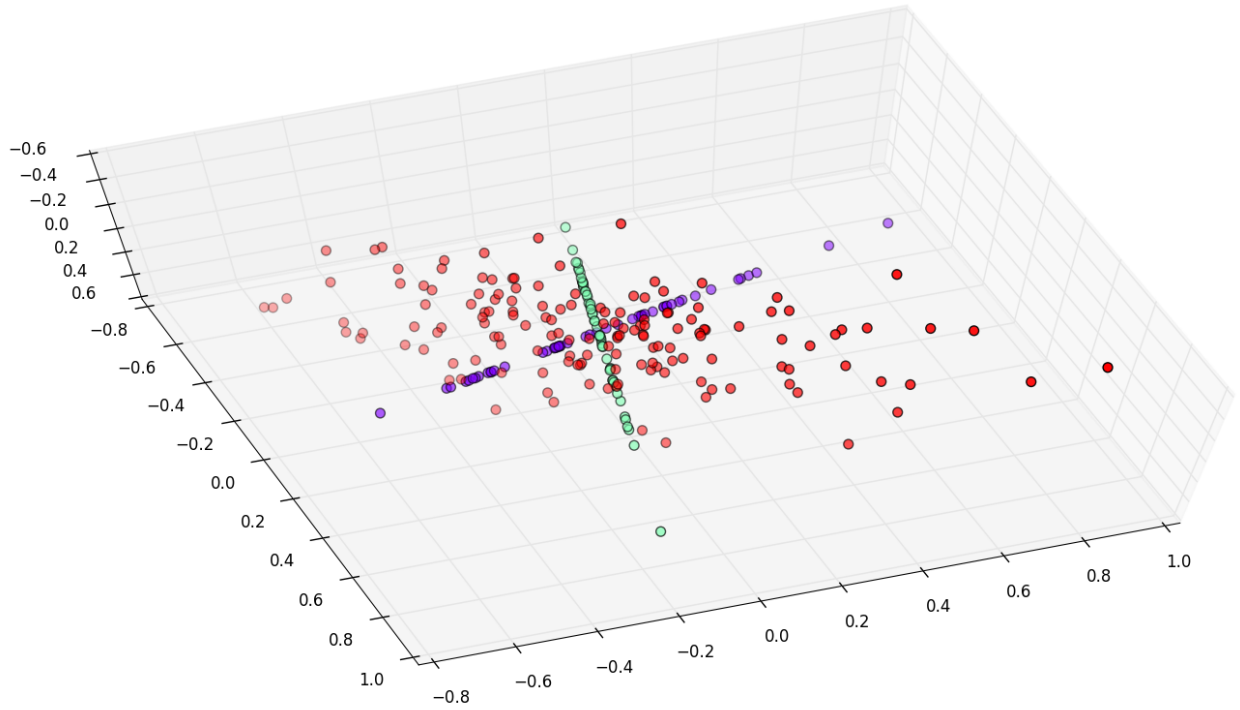

Figure 1: Three subspaces in $\mathbb{R}^3$: the first and second group of 50 data points belong to two one-dimensional subspaces, and the remaining 150 data points belong to a two-dimensional subspace.

2 Subspace Clustering

2.1 Overview

Subspace Clustering consists in finding low-dimensional representation of high-dimensional data. For example, given data points drawn from a three-dimensional space, we recover their subspace structures via subspace clustering as depicted in Figure 1.

The subspace clustering problem can be framed mathematically as follows.

Problem 1 (Subspace clustering)**.** *Let $X \in \mathbb{R}^{n \times N}$ be a real-valued matrix whose columns are N data points drawn from a union of M subspaces S_j of $\mathbb{R}^n$, of dimensions $d_j < \min\{n, N\}$, for $j = 1, \ldots, M$. Given X, the subspace clustering seeks to find: a) the subspace to which each data point belongs to. That is, the segmentation of the data points; b) the number of subspaces and their intrinsic dimensions.*

Note that finding the respective basis of subspaces can be obtained as a by-product of solving the subspace clustering problem. Recent approaches to solving problem 1 are based on the *self-expressiveness model* [17, 11], which states that each data point in a union of subspaces can be expressed as a linear or an affine combination of other points. When data points lie in a linear subspace S_j of dimension d_j, each data point can be written as a linear combination of d_j other points from S_j. Similarly, when data points lie in an affine subspace S_j of dimension d_j, each data point can be written as an affine combination of $d_j + 1$ other points from S_j. The combination can be represented using matrix notation by $X = XZ$, where X denotes the *data matrix* and Z the *matrix of coefficients*.

Let us first consider the problem 1 in the case of linear subspaces. Although this combination is in general not unique, a *sparse representation* of a data point ideally corresponds to a combination of a few points belonging to its own subspace [6]. By solving a global sparse optimization problem, and obtaining

the matrix of coefficients, the spectral clustering technique is then applied to infer the clustering of data points, whose theoretical bound can be guaranteed by the normalized cut [16]. In addition, the number of the nonzero elements in the sparse representation of a data point ideally corresponds to the dimension of its subspace. As for the number of subspaces, it can be found either by analyzing the eigenspectrum of the Laplacian matrix of the graph constructed by Z [21], or by model selection techniques [3]. Finally, the respective basis of subspaces can be readily obtained. Therefore, with the sparse representation framework based on self-expressiveness, we can solve thoroughly the subspace clustering problem 1. Compared to local spectral clustering-based algorithms, we can furthermore overcome challenges such as dealing with points near the intersection of subspaces.

When it comes to the case of affine subspaces, the property that a data point can be represented sparsely by a few points from its own subspace becomes more subtle. [12] introduces situations where the property fails to be met when a grouping effect for interior points is present. Nonetheless, our experiments show that the sparse self-expressiveness in the case of affine subspaces is satisfied in our context.

2.2 Structured sparse subspace clustering ($\mathrm{S}^3\mathrm{C}$)

State-of-the-art subspace clustering methods based on the self-expressiveness model follow a two-step approach to solve the subspace clustering problem 1: learning an affinity matrix from data in the first step and segmenting data based on the affinity matrix in the second step. This two-step approach is generally suboptimal since it does not consider the interdependence of the affinity with the segmentation. In [11], Li et al. seek to combine these two steps by building a unified optimization framework, dubbed *Structured Sparse Subspace Clustering* (SSSC or $\mathrm{S}^3\mathrm{C}$) in which obtaining affinity matrix and segmentation of data are merged together. More specifically, this approach considers the $\mathrm{S}^3\mathrm{C}$ problem as follows:

$$\begin{aligned} & l \min_{Z,E,Q} \|Z\|_{1,Q} + \lambda \|E\|_{\ell} \\ \text{s.t. } & X = XZ + E,\ \mathrm{diag}(Z) = \mathbf{0},\ Q \in \mathcal{Q}. \end{aligned} \tag{1}$$

where X is the $n \times N$ data matrix and Z the coefficient matrix. The norm $\|\cdot\|_{\ell}$ on the error term E depends upon the prior knowledge about the pattern of noise. $Q = [\mathbf{q_1}, \cdots, \mathbf{q_M}]$ refers to an $N \times M$ *binary segmentation matrix* indicating the membership of each data point to each subspace. That is, $q_{i,j} = 1$ if the i-th column of X lies in subspace S_j and $q_{i,j} = 0$ otherwise. The space of segmentation matrices is defined as $\mathcal{Q} = \{Q \in \{0,1\}^{N \times M} : Q\mathbf{1} = \mathbf{1} \text{ and } \mathrm{rank}(Q) = M\}$. λ and α are two trade-off parameters. The *subspace structured* ℓ_1 norm of Z is defined as follows:

$$\begin{aligned} rCl\, \|Z\|_{1,Q} &\doteq \|Z\|_1 + \alpha \|\Theta \odot Z\|_1 \\ &= \sum_{i,j} |Z_{i,j}| \left(1 + \frac{\alpha}{2} \left\| \boldsymbol{q}^{(i)} - \boldsymbol{q}^{(j)} \right\|^2 \right) \end{aligned} \tag{2}$$

where $\Theta_{i,j} = \frac{1}{2} \left\| \boldsymbol{q}^{(i)} - \boldsymbol{q}^{(j)} \right\|^2$ with $\boldsymbol{q}^{(i)}$ and $\boldsymbol{q}^{(i)}$ the i-th and j-th row of matrix Q, respectively.

The optimization problem (1) ensures a sparse representation of data points by means of the ℓ_1 norm. The constraint $X = XZ + E$ seeks to capture the self-expressiveness while $\mathrm{diag}(Z) = \mathbf{0}$ prevents data points representing themselves. The introduction of (2) merges the affinity and segmentation together. By formulating the subspace clustering problem in this way, $\mathrm{S}^3\mathrm{C}$ is able to yield more coherent clustering results. We build our unified framework for video co-segmentation based on $\mathrm{S}^3\mathrm{C}$.

3 A unified framework for video co-segmentation

In this section, we introduce our unified framework for video co-segmentation with *appearance* and *motion features*. Appearance feature and motion feature are both crucial in object segmentation and recognition.

On one hand, the appearance feature captures color or texture information of objects. It emphasizes the salience of an object by which it stands out relative to its neighbors. On the other hand, the motion feature shows how objects move. By exploiting together the appearance and motion information, we simulate thus how human beings recognize and differentiate objects. We given an overview of our framework as follows:

First, we build temporal superpixels in the preprocessing stage in order to reduce computational costs. We extract motion and appearance features from those superpixels. Furthermore, we improve the detectability of motion features by adding an *extra signature* to the conventional motion trajectory matrix. Then, we formulate the unified optimization framework with the additional affine subspace constraint, and solve our global optimization problem efficiently using the Alternating Direction Method of Multipliers (ADMM) method. Finally, we construct a unified affinity matrix from the optimal solution of the optimization problem, and obtain segmentation results by applying spectral clustering to the affinity matrix.

3.1 Preprocessing: Temporal superpixels

In order to reduce the computational complexity, we represent videos using Temporal SuperPixels (TSPs) [4] which help build superpixel-wise spatio-temporal correspondences. TSPs are different from common over-segmentation techniques because those superpixel representations maintain the most salient features of an image. In addition, we integrate SIFT flow [13, 22] into the TSP framework so as to improve the robustness of correspondences. We then extract motion and appearance features at the level of TSPs.

3.2 Appearance features

We choose the appearance feature to be 10D color features in HSV color space. Gaussian mixture model and discretization are then applied to the color features to obtain the color histogram [24].

3.3 Motion feature using an extra-dimensional signature

In order to capture motion information of objects, we need a function called a *camera model* to map the three-dimensional world onto a two-dimensional *image plane*. In the case of imaging objects far away from a camera, small differences in depth, or the *perspective effect*, becomes less apparent [15]. The perspective projection can thus be approximated by an *affine projection* with the *affine camera model*. Thus we construct the motion feature under the affine projection model to fit the affine subspace clustering method.

Traditionally [20], motion feature matrix is represented by (3).

$$lM = \begin{bmatrix} x_{11} & \dots & x_{1N} \\ \vdots & & \vdots \\ x_{F1} & \dots & x_{FN} \end{bmatrix}_{2F \times N} \tag{3}$$

where $\{x_{\mathrm{fj}} \in \mathbb{R}^2\}_{j=1,\dots,N}^{f=1,\dots,F}$ denotes the 2D projection of N 3D points on rigidly moving objects onto F frames of a moving camera. In our case of using superpixels, x_{fj} refers to the center of the intersection of j-th TSP and f-th frame. However, in real-world problems, there often exist missing trajectories caused by tracked points moving into or out of frames, or by occlusion. Conventionally, nearest motion values in the same data column are padded to the missing entries [22]. While this simple treatment sometimes results in acceptable segmentation, generally it reduces the detectability of motion trajectories, making subspace clustering method unable to perform motion segmentation. To overcome this limitation, in addition to padding motion values, we propose a new way to form motion feature matrix by adding an *extra-dimension signature* so that $\{x_{\mathrm{fj}} \in \mathbb{R}^3\}_{j=1,\dots,\mathrm{N}}^{f=1,\dots,F}$ with $x_{\mathrm{fj}} = (x\ y\ s)^\top$. s denotes the signature added to the original motion trajectories. Different values of s indicate different *status* of motion trajectories, such as missing trajectories caused by newly appearing points or by dead points, which can be readily inferred from indices of TSPs . Newly appearing TSPs have indices that do not exist in previous frames while indices assigned to dead TSPs disappear starting from certain frame. The values of the signature, $s_i,\ i = 1, 2, 3$ are specified

(a) Original frame

(b) With signature

(c) Without signature

Figure 2: Comparison of motion feature with and without signature.

so as to encode the status of motion trajectories and sufficiently distinguish one from another. Here, we set $s_i,\ i = 1, 2, 3$ to be $1, 3, 5$, respectively, as shown in Table 1.

Value	Status of motion trajectories
$s_1 = 1$	Original complete trajectories
$s_2 = 3$	Missing trajectories caused by new points
$s_3 = 5$	Missing trajectories caused by dead points

Table 1: Meaning of signature s to status of motion trajectories: Depending on newly appearing or disappearing feature points, we encode the event using an extra dummy dimension to get better and more consistent subspace clustering results.

By marking the difference by means of the signature, motion features are more detectable in the presence of missing trajectories, which is demonstrated experimentally in Figure 2. The subspace clustering method succeeded in segmenting motions with signature (2b) while it fails in the absence of signature (2c).

3.4 Unified optimization framework

We propose the following unified optimization framework (4). In order to make it work on affine subspaces, we add an affine subspace constraint $Z_k^\top \mathbf{1} = \mathbf{1}$ to the original S^{3}C framework (1), and formulate the optimization problem to solve as follows:

$$C \min_{\substack{Z_1,Z_2,\ldots,Z_K \\ E_1,E_2,\ldots,E_K \\ Q_1,Q_2,\ldots,Q_K}} \sum_{k=1}^{K} (\|Z_k\|_{1,Q_k} + \lambda_k \|E_k\|_1) + \beta \|Z\|_{2,1}$$
$$s.t. \quad X_k = X_k Z_k + E_k,\ \mathrm{diag}(Z_k) = \mathbf{0},$$
$$Z_k^\top \mathbf{1} = \mathbf{1},\ Q_k \in \mathcal{Q}$$

where

$$\begin{aligned} Z &= \begin{pmatrix} \mathrm{vec}(Z_1^\top) & \mathrm{vec}(Z_2^\top) & \cdots & \mathrm{vec}(Z_K^\top) \end{pmatrix}^\top \\ &= \begin{bmatrix} (Z_1)_{11} & (Z_1)_{12} & \cdots & (Z_1)_{\mathrm{NN}} \\ (Z_2)_{11} & (Z_2)_{12} & \cdots & (Z_2)_{\mathrm{NN}} \\ \vdots & \vdots & & \vdots \\ (Z_K)_{11} & (Z_K)_{12} & \cdots & (Z_K)_{\mathrm{NN}} \end{bmatrix}_{K \times N^2} \end{aligned}$$

and $\|Z_k\|_{1,Q_k} = \|Z_k\|_1 + \alpha \|\Theta_k \odot Z_k\|_1$. $\mathbf{0}$ and $\mathbf{1}$ are two vectors of appropriate dimensions filled with 0 and 1, respectively. X_k refers to different feature data points. Here, with $K = 2$, X_1 denotes appearance features and X_2 denotes motion features. Z_k is the coefficient matrix and E a matrix of errors. The constraint $Z_k^\top \mathbf{1} = \mathbf{1}$ generalizes the framework to affine subspaces by writing each data point as an affine combination of a few other points. Since the linear subspace is a special case of the affine subspace, our framework works effectively also on data points lying in linear subspaces.

Algorithm 1 Unified framework for solving problem (4) with affine subspace constraint

Input: Data matrix X, α, β, λ, ρ, and the number of subspaces.

1. **Initialize:** $\Theta_0 = \mathbf{0}$, $\epsilon = 10^{-6}$
2. **while** not converged **do**
3. Given Θ_T, initialize $E = \mathbf{0}$, $C = Z = J = \mathbf{0}$, $Y^{(1)} = \mathbf{0}$, $Y^{(2)} = Y^{(3)} = \mathbf{0}$, $Y^{(4)} = \mathbf{0}$
4. **while** not converged **do**
5. Update J_t, C_t, E_t, and Z_t according to (4) (4) (5) (7);
6. Update $Y_t^{(1)}$, $Y_t^{(2)}$, $Y_t^{(3)}$, $Y_t^{(4)}$ according to (7);
7. Update $\mu_{t+1} \leftarrow \rho\mu_t$;
8. Check the condition $\|X - XC - E\|_\infty < \epsilon$; if not converged, then set $t \leftarrow t+1$; if converged, obtain Z
9. **end while**
10. Given Z, construct the affinity graph G given by $\mathbf{W} = |\mathbf{Z}| + |\mathbf{Z}|^{\mathrm{T}}$.
11. Apply the spectral clustering technique to the affinity graph with the input number of subspaces as k and obtain Θ_{T+1}.
12. Check the convergence condition $\|\Theta_{T+1} - \Theta_T\|_\infty < 1$; if not converged, $T \leftarrow T+1$
13. **end while**

Output: Coefficient matrix Z and the label l indicating subspaces where data points lie in

$\|Z\|_{2,1}$ is a penalty term which encourages a unified affinity matrix inferred from appearance and motion features since the $\ell_{2,1}$ norm can induce column sparsity of the matrix [22].

To solve the optimization problem (4), we notice that it is equivalent to the following problem (4) which can be solved more efficiently.

$$
\begin{aligned}
C \min_{\substack{Z_1,Z_2,\ldots,Z_K \\ E_1,E_2,\ldots,E_K \\ Q_1,Q_2,\ldots,Q_K}} & \sum_{k=1}^{K}(\|J_k\|_{1,Q_k} + \lambda_k\|E_k\|_1) + \beta\|Z\|_{2,1} \\
& s.t. \quad X_k = X_kC_k + E_k,\ J_k = Z_k, \\
& C_k = J_k - \mathrm{diag}(J_k),\ C_k^\top \mathbf{1} = \mathbf{1},\ Q_k \in \mathcal{Q}
\end{aligned}
$$

We can solve this problem using the Alternating Direction Method of Multipliers (ADMM) method [2], which is a powerful yet simple algorithm well suited to problems arising in applied statistics and machine learning. The augmented Lagrangian is given by:

$$
\begin{aligned}
& l\mathcal{L}(J_k, C_k, E_k, Z, Y_k^{(1)}, Y_k^{(2)}, Y_k^{(3)}, Y_k^{(4)}; k = 1, \ldots, K) \\
= & \sum_{k=1}^{K} \|J_k\|_{1,Q_k} + \lambda_k\|E_k\|_1 + \langle Y_k^{(1)}, X_k - X_kC_k - E_k\rangle \\
& + \langle Y_k^{(2)}, C_k - J_k + \mathrm{diag}(J_k)\rangle + \langle Y_k^{(3)}, J_k - Z_k\rangle \\
& + \langle Y_k^{(4)}, C_k^\top \mathbf{1} - \mathbf{1}\rangle + \frac{\mu}{2}(\|X_k - X_kC_k - E_k\|_F^2 \\
& + \|C_k - J_k + \mathrm{diag}(J_k)\|_F^2 + \|J_k - Z_k\|_F^2 \\
& + \|C_k^\top \mathbf{1} - \mathbf{1}\|_F^2) + \beta\|Z\|_{2,1},
\end{aligned}
$$

where $Y_k^{(1)}, Y_k^{(2)}, Y_k^{(3)}, Y_k^{(4)}$ are Lagrange multipliers. ADMM allows us to update variables alternatively. In addition, since $J_k, C_k, E_k, Y_k^{(1)}, Y_k^{(2)}, Y_k^{(3)}, Y_k^{(4)}$ with $k = 1, \ldots, K$ are separable in the optimization problem, we can update them separately given k (using OpenMP for multi-cores processors).

Update J_k. We update J_k by solving the following problem:

$$\begin{aligned} lllJ_{k,t+1} &= \arg\min_{J_k} \frac{1}{\mu_t}\|J_k\|_{1,Q_k} + \frac{1}{2}\|J_k - \mathrm{diag}(J_k) - U_{k,t}\|_F^2 \\ &+ \frac{1}{2}\|J_k - Z_{k,t} + Y_{k,t}^{(3)}/\mu_t\|_F^2, \end{aligned}$$

where $\|J_k\|_{1,Q_k} = \|J_k\|_1 + \alpha\|\Theta_k \odot J_k\|_1$ and $U_{k,t} = C_{k,t} + \frac{1}{\mu_t}Y_{k,t}^{(2)}$.
The closed-form solution for J_k is given by

$$lJ_{k,t+1}^{i,j} = \begin{cases} \frac{1}{2}S_{\frac{1}{\mu_t}(1+\alpha\Theta_k^{i,j})}\left(U_t^{i,j} + Z_{k,t}^{i,j} - \frac{Y_{k,t}^{(3),i,j}}{\mu_t}\right) & i \neq j, \\ S_{\frac{1}{\mu_t}(1+\alpha\Theta_k^{i,j})}\left(Z_{k,t}^{i,j} - \frac{Y_{k,t}^{(3),i,j}}{\mu_t}\right) & i = j. \end{cases}$$

where $U_{k,t} = C_{k,t} + \frac{1}{\mu}Y_{k,t}^{(2)}$. $S_\tau(\cdot)$ is the shrinkage thresholding operator acting on every element of a matrix. It is defined as: $S_\tau(\nu) = (|\nu| - \tau)_+ \,\mathrm{sgn}(\nu)$ with the operator $(\nu)_+$ returning ν if it is non-negative and returning zero otherwise.
Update C_k. We update C_k by solving the following problem:

$$\begin{aligned} lllC_{k,t+1} &= \arg\min_{C_k} \langle Y_{k,t}^{(1)}, X_k - X_kC_k - E_{k,t}\rangle \\ &+ \langle Y_k^{(2)}, C_k - J_{k,t+1} + \mathrm{diag}(J_{k,t+1})\rangle \\ &+ \langle Y_{k,t}^{(4)}, C_k^\top \mathbf{1} - \mathbf{1}\rangle + \frac{\mu_t}{2}(\|C_k^\top \mathbf{1} - \mathbf{1}\|_F^2 \\ &+ \|C_k - J_{k,t+1} + \mathrm{diag}(J_{k,t+1})\|_F^2 \\ &+ \|X_k - X_kC_k - E_{k,t}\|_F^2) \end{aligned}$$

The solution is given by

$$\begin{aligned} lllC_{k,t+1} &= (X_k^\top X_k + \mathrm{I} + \mathbf{1}\mathbf{1}^\top)^{-1}\left[X_k^\top\left(X_k - E_{k,t} + \frac{1}{\mu_t}Y_{k,t}^{(1)}\right)\right. \\ &+ J_{k,t+1} - \mathrm{diag}(J_{k,t+1}) - \frac{1}{\mu_t}Y_{k,t}^{(2)} \\ &- \left.\frac{1}{\mu_t}(\mathbf{1}Y_{k,t}^{(4)\top}) + \mathbf{1}\mathbf{1}^\top\right] \end{aligned}$$

Update E_k. We update E_k by solving the following problem:

$$E_{k,t+1} = \arg\min_{E_k} \frac{\lambda_k}{\mu_t}\|E_k\|_1 + \frac{1}{2}\|E_k - V_{k,t}\|_F^2, \tag{4}$$

where $V_{k,t} = X_k - X_kC_{k,t+1} + \frac{1}{\mu_t}Y_{k,t}^{(1)}$.
Then the solution is given by

$$E_{k,t+1} = \quad S_{\frac{\lambda}{\mu_t}}(V_{k,t}) \tag{5}$$

where $V_{k,t} = X_k - X_kC_{k,t+1} + \frac{1}{\mu_t}Y_{k,t}^{(1)}$.
Update Z. We solve Z by solving the following problem:

$$Z_{t+1} = \quad \arg\min_{Z} \frac{\beta}{\mu_t}\|Z\|_{2,1} + \frac{1}{2}\|Z - P_t\|_F^2, \tag{6}$$

Algorithm 2 $S^3C - AM$

Input: Feature matrices X_k, $k = 1, \ldots, K$, α, β, λ, ρ, and the number of subspaces.

1. Obtain Z_k, $k = 1, \ldots, K$ via Algorithm 1;
2. Form the affinity matrix S according to (7);
3. Apply the spectral clustering technique to the affinity matrix S with the input number of subspaces;

Output: The label l indicating subspaces where TSPs lie in

where

$$\begin{aligned} P_t &= \left[\begin{array}{cccc} \mathrm{vec}(P_1^\top) & \mathrm{vec}(P_2^\top) & \cdots & \mathrm{vec}(P_K^\top) \end{array} \right]^\top \\ &= \left[\begin{array}{cccc} (P_1)_{11} & (P_1)_{12} & \cdots & (P_1)_{\mathrm{NN}} \\ (P_2)_{11} & (P_2)_{12} & \cdots & (P_2)_{\mathrm{NN}} \\ \vdots & \vdots & & \vdots \\ (P_K)_{11} & (P_K)_{12} & \cdots & (P_K)_{\mathrm{NN}} \end{array} \right]_{K \times N^2} \end{aligned}$$

with $P_{k,t} = J_{k,t+1} + \frac{Y_{k,t}^{(3)}}{\mu_t}$

Then the problem can be solved via Lemma 4.1 of [14] as follows:

$$l[Z_{t+1}]_{:,i} = \left\{ \begin{array}{cc} \frac{\|[P_t]_{:,i}\|_2 - \frac{\beta}{\mu_t}}{\|[P_t]_{:,i}\|_2}[P_t]_{:,i} & \text{if } \|[P_t]_{:,i}\|_2 > \frac{\beta}{\mu_t}, \\ 0 & \text{otherwise}. \end{array} \right.$$

where $[Z_{t+1}]_{:,i}$ and $[P_t]_{:,i}$ refer to the i^{th} column of respective matrix. $Z_{k,t+1}$ can then be obtained from Z_{t+1}.

Update $Y_k^{(1)}, Y_k^{(2)}, Y_k^{(3)}, Y_k^{(4)}$. We update the Lagrange multipliers by a simple gradient ascent step.

$$\begin{aligned} lclY_{k,t+1}^{(1)} &= Y_{k,t}^{(1)} + \mu_t(X_k - X_k C_{k,t+1} - E_{k,t+1}) \\ Y_{k,t+1}^{(2)} &= Y_{k,t}^{(2)} + \mu_t(C_{k,t+1} - J_{k,t+1} + \mathrm{diag}(J_{k,t+1})) \\ Y_{k,t+1}^{(3)} &= Y_{k,t}^{(3)} + \mu_t(J_{k,t+1} - Z_{k,t+1}) \\ Y_{k,t+1}^{(4)} &= Y_{k,t}^{(4)} + \mu_t(C_{k,t+1}^\top \mathbf{1} - \mathbf{1}) \end{aligned}$$

Finally, we update μ by $\mu_{t+1} \leftarrow \rho\mu_t$. The unified framework is summarized in Algorithm 1. For the details of the derivation, we refer the readers to [2].

3.5 Co-segmentation algorithm

Finally, in order to combine appearance and motion features, we construct the affinity matrix S given by

$$l(S)_{ij} = \frac{1}{2}\left(\sqrt{\sum_{k=1}^{K} (Z_k)_{ij}^2} + \sqrt{\sum_{k=1}^{K} (Z_k)_{ji}^2} \right) \tag{7}$$

Then, we apply the spectral clustering technique [21] to the affinity matrix S to obtain segmentation results in the form of an array of label l indicating corresponding subspace of each data point. For clarity, we summarize the $S^3C - AM$ algorithm in Algorithm 2. The overall processing is summarized in Figure 3.

4 Experiments

In this section, we conduct experiments on the dataset Multi-Object Video Co-Segmentation (MOViCS)[1] proposed by [5] to evaluate the effectiveness of the $S^3C - AM$ framework. This dataset is composed of four

[1]The MOViCS dataset can be found at `http://www.d2.mpi-inf.mpg.de/datasets`

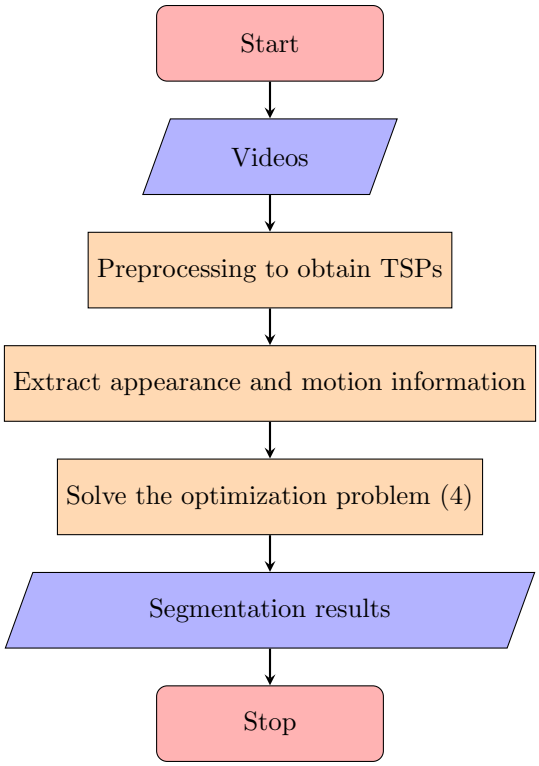

Figure 3: Flow chart of the overall processing

sets of real videos, each of which contains different combinations of animal objects: a) chicken-turtle; b) giraffe-elephant; c) zebra-lion; d) tiger. The objects in the sets of videos present various appearance and motion patterns, making the co-segmentation task more challenging. Some animals may move inseparably together such as in the *chicken on turtle* video. Natural coloring of some other animals, including elephant or tiger, provides camouflage so that they blend in perfectly with the surroundings. Several other difficulties include different lighting conditions, and motion blur. Despite those challenges, our experimental results show that our framework achieves the highest overall performance, and outperforms baseline algorithms in two sets of videos and yields comparable results on two remaining video sets.

4.1 Parameter setting

In $\mathrm{S^3C-AM}$, the parameters λ_k and α need to be set properly as in [6]. Parameters ρ and β are set to be 1.2 and 1×10^{-5}, respectively, which work well in practice.

4.2 Evaluation metric

In order to evaluate the effectiveness of our framework, we use the following metric:

$$\text{Score} = \frac{1}{C} \sum_j \max_i \frac{S_i \cap G_j}{S_i \cup G_j}, \tag{8}$$

where G_j and S_i refer to sets of segments belonging to j^{th} class obtained by our co-segmentation algorithm, and i^{th} class in ground truth, respectively. C is the number of classes in ground truth. By computing the average accuracy that a set of segments matches best each object in video frames, the *intersection-over-union metric* (8) is able to reflect the overall performance of a video co-segmentation method.

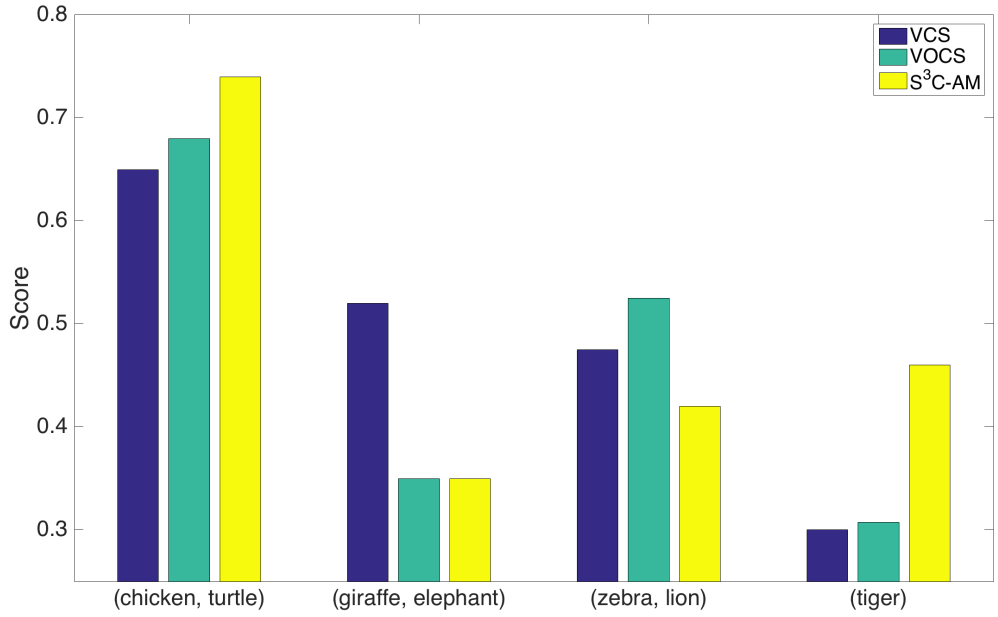

Figure 4: Comparison of co-segmentation accuracies between our method (S^3C-AM), multi-class video co-segmentation (VCS) and video object co-segmentation (VOCS) on MOViCS dataset.

4.3 Results

We conduct experiments on all video sets in the dataset and compare our framework to other state-of-the-art methods [5] and [22]. For the purpose of clarity, we denote the *multi-class video co-segmentation* framework of [5] by VCS and the *video object co-segmentation* framework of [22] by VOCS. We refer to our framework as S^3C-AM. For both VCS and VOCS, we use their publicly available accuracies measured also by the metric (8) on the MOViCS dataset.

Figure 4 shows the comparison between our framework and the baselines. Achieving an overall performance of 49.25% by our method, we outperform VCS (48.75%) and VOCS (46.56%). For two out of four video sets, i.e., chicken-turtle, and tiger videos, our framework produces better results, particularly for the tiger video set where undistinguishable appearance of the tiger with respect to its surroundings is essentially present. For the giraffe-elephant video set, VCS outperforms both S^3C-AM and VOCS while for the zebra-lion video set our method produces lower accuracy compared to two baselines. We will give full discussion in the later subsection.

In addition, we compare computational cost between VCS, VOCS, and our framework. We evaluate our framework under the following computer configuration: 2.7 GHz Intel Core i5, with 8 GB RAM and OS platform, as shown in Table 2. Our preprocessing performs superpixel over-segmentation which reduces greatly the number of data points involved in clustering. While [22] adopts also TSP-based preprocessing, our method is able to achieve reliable co-segmentation results using less TSPs, leading to faster preprocessing. For three out of four video sets our clustering routine runs much faster than VCS. Due to our iterative clustering algorithm intended for more coherent segmentation, it is expected that our method needs more time in clustering part than VOCS.

Table 3 shows visualization of our results. Every block corresponds to frame-result pairs of a video set. First column in each block is a frame of the video, whereas second column shows the results of our method

	Preprocessing			Clustering		
	VCS	VOCS	S^{3}C-AM	VCS	VOCS	S^{3}C-AM
(chicken, turtle)	1h 40m	1h 10m	31m	1h 12m	1h 04m 13s	1h 18m
(giraffe, elephant)	1h 33m	54m	22m	1h 22m	13m 15s	35m
(zebra, lion)	3h 19m	2h 14m	56m	3h 13m	1h 23m 20s	2h 11m
(tiger)	1h 11m	48m	21m	59m	30m 18s	46m

Table 2: Comparison of computation cost between VCS, VOCS, and our method S^{3}C-AM. The runtime data of [5] and [22] are obtained from their supplementary material.

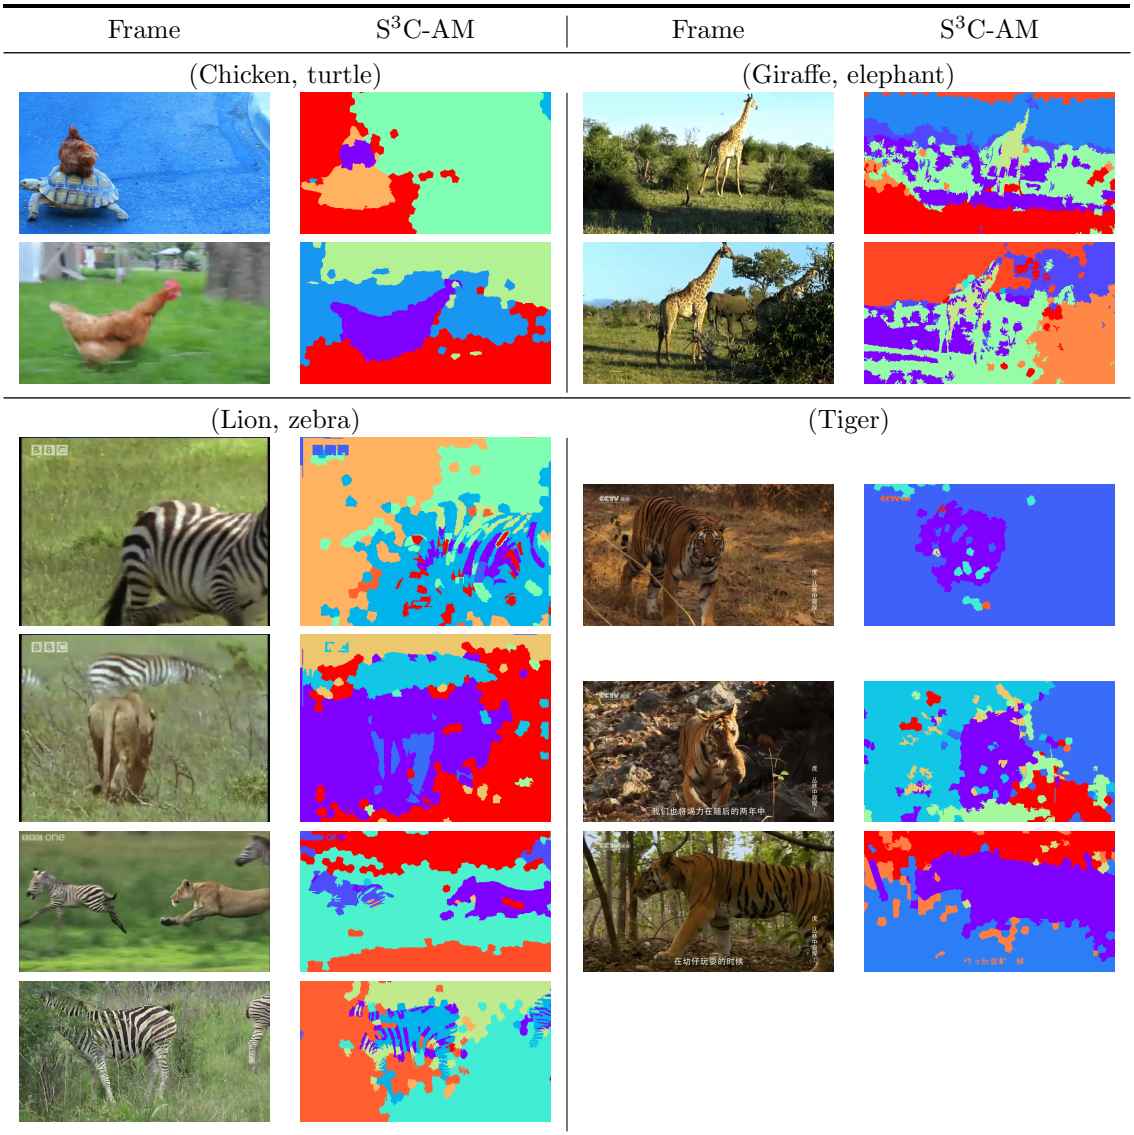

Table 3: Example of results of co-segmentation for all video sets in MOViCS dataset.

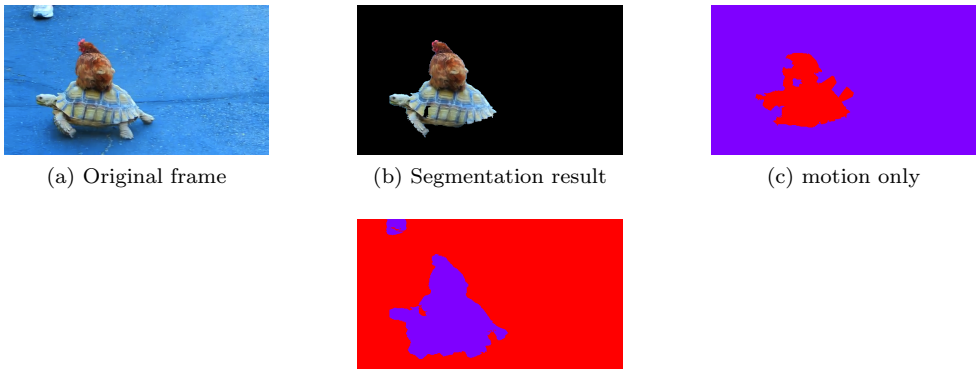

(a) Original frame (b) Segmentation result (c) motion only

(d) appearance only

Figure 5: Segmentation results on the *chicken on turtle* video sequences

4.4 Discussion

4.4.1 Appearance and motion features

Using both appearance and motion information enables our method to better segment objects. For illustration, we compare segmentation of (chicken on turtle) video only with the motion feature (5c) to that solely with the appearance features (5d). It can be noticed that, with only appearance features, objects with similar colors in the background may be segmented to the same class of the foreground. On the other hand, segmentation without appearance information may result in incomplete segmentation. The combination of appearance and motion features gives a satisfactory segmentation (5b) even in the presence of noise and blur found in the original videos.

4.4.2 Affine subspace constraint

As stated before, we added an affine subspace constraint $Z_k^\top \mathbf{1} = \mathbf{1}$ to the original S^{3}C framework (1), which we believe makes ours yield better results on affine subspaces. In order to highlight the role of the affine subspace constraint with respect to the motion feature, we compare our framework with S^{3}C on the Hopkins 155 motion segmentation dataset [18] that allows us to analyze the ability of both frameworks to segment video sequences into multiple spatio-temporal regions corresponding to the different motions in the scene without interference of the appearance feature. In the dataset, feature points are already tracked so that no preprocessing of temporal superpixels is needed. Since our framework is compatible with a single feature by removing the penalty term $\|Z\|_{2,1}$, we set β to be 0 to ensure an equivalent parameter setting. For both frameworks, we feed motion features without the extra-dimensional signature.

The evaluation metric of clustering is given by

$$\text{error}(l, \hat{l}) = 1 - \max_{\pi} \frac{1}{N} \sum_{i=1}^{N} 1_{\{\pi(\hat{l})=l\}}, \tag{9}$$

where $l, \hat{l} \in \{1, \ldots, M\}^N$ are original and estimated labels of tracked points in the M subspaces. As the labels assigned to the points by subspace clustering algorithms may be a permutation of the original ones, we find the most coherent labels by maximizing the second term with respect to all permutations $\pi : \{1, \ldots, M\}^N \to \{1, \ldots, M\}^N$.

Experimental results of the comparison are presented in Table 4. It is shown that with the additional affine subspace constraint, our framework is able to achieve higher accuracy when segmenting rigid motions lying in affine subspaces. The affine subspace constraint, which improves the performance of motion segmentation, as well as the combination of motion and appearance features, ensure the performance of our framework.

No. motions	2		3	
Error(%)	Ave.	Med.	Ave.	Med.
S^3C [11]	1.94	**0**	4.92	0.89
S^3C-AM	**1.53**	**0**	**4.40**	**0.56**

Table 4: Comparison of motion segmentation errors on Hopkins 155 Database between S^3C (without the affine subspace constraint) and S^3C-AM (with the affine subspace constraint). The best results are in bold font.

4.4.3 Problems caused by missing trajectories

Our framework exploits the superpixel over-segmentation. Thus, the performance of our method depends on results of the TSP preprocessing. While TSP method outperforms other state-of-the-art approaches and provides reliable segmentation results, it produces dead TSPs and new TSPs due to occlusion or tracked points moving into or out of frames. Consequently, TSPs can only survive during a limited number of frames. Such phenomenon would greatly reduce the inter-frame temporal extent of TSPs, resulting in many missing trajectories, especially when objects move rapidly or encounter frequent occlusion caused by themselves or others. This is exactly the case in (giraffe, elephant) and (zebra, lion) video sets. Even though we propose a novel and effective method to deal with missing trajectories through an extra-dimensional signature, we lose too many trajectories to have the signature compensate for the loss of information. Since our method emphasizes essentially the coherence of both appearance and motion features, a flaw in motion features would finally result in failure of the unified framework. Therefore, due to lack of sufficient motion features related to (giraffe, elephant) and (lion, zebra) videos, our method fails to fully discover objects in video sets.

4.4.4 Multi-class segmentation vs. co-segmentation

As we understand, multi-class segmentation and co-segmentation both refer to the process of assigning a label to every pixel in videos and partitioning them into meaningful groups. They are closely related concepts while they differ in the way to provide additional information in the absence of supervisory information. Co-segmentation seeks to segment an object in video groups where it all appears. As such, the presence of objects shared among videos provides more information and encourages co-segmentation. For example, given a video of a chicken on a turtle and the other video of only a chicken, our framework is able to extract the common object, i.e., the chicken and thus obtain a refined result as shown in Table 5. On the other hand, multi-class segmentation does not enforce such setting. It usually discovers object segmentation based on a single video where salient objects are present. The underlying assumption of multi-class segmentation and co-segmentation about the structure of videos may coincide. It is also possible for co-segmentation to assume existence of several classes, or only background/foreground binary classes. Our framework is more like multi-class co-segmentation in the sense that it assumes multi-class structure of videos and yields multi-class segmentation results. It is also possible for our framework to yield binary segmentation if videos are less structured as in the case of (chicken, turtle). In general cases, further processing is needed to highlight foreground/background, such as user guidance.

4.4.5 Robustness under the condition of heavy noise

Finally, we tested the robustness of our framework under the condition of heavy noise. While for most videos there is no concern about resolution and noise, in some circumstances it is inevitable to encounter video quality issues. For example, scenes shot at night under poor lighting conditions are prone to contain heavy noise due to high ISO sensitivity. We added gaussian-distributed additive noise with variance of 0.25, or 5% amount of salt-and-pepper noise to the original video. It is shown that our framework is effective even with heavy data corruption as illustrated in Figure 6.

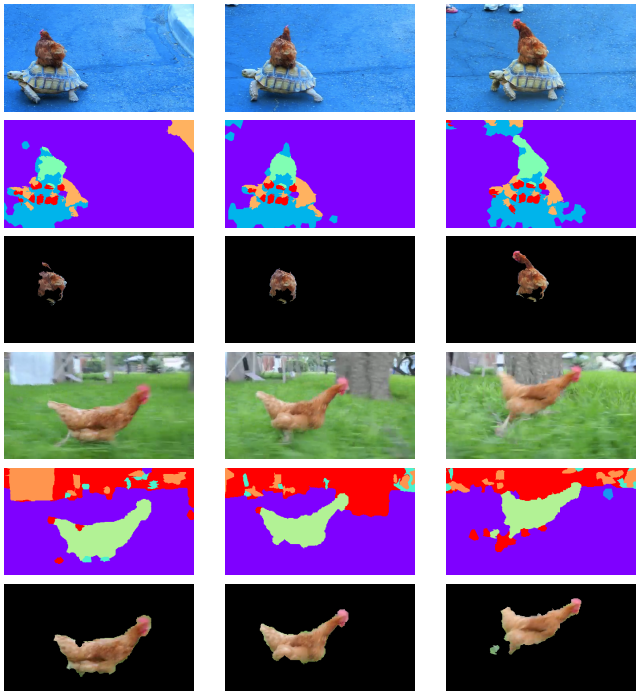

Table 5: Results of co-segmentation on videos *chicken on turtle* and *chicken*. Columns 1, 2, and 3 correspond to frames 1, 9, and 17. Rows 1 and 4 correspond to the original video frames, rows 2 and 5 are multi-class segmentation, rows 3 and 6 correspond to the extracted common objects.

5 Conclusion and perspectives

We described a novel unified framework for video co-segmentation problem. By adding an extra-dimensional signature to motion trajectories, we improve the detectability of motion features in the subspace clustering framework. Moreover, we reformulated the S^3C with the additional affine subspace constraint to successfully segment data points drawn from a union of affine subspaces. We then formulated the unified optimization framework by integrating the newly proposed subspace clustering algorithm and applied it to appearance and motion features. Experiments on the MOViCS benchmark demonstrated the effectiveness of our method and its superiority over baselines. In future work, we would like to extend our approach by generalizing the framework to objects with articulated motion. In addition, we plan also exploit a distributed optimization based on ADMM to improve the running times. We concisely hint at those extensions as follows:

5.1 Object with articulated motion

As shown by the experiments, our approach segments effectively salient objects from videos by jointly making full use of the appearance and motion information. As for the motion information, it is assumed that each object has a rigid-body motion which forms an affine subspace. Though this assumption works well in general, it is difficult to handle objects with articulated motion. To generalize our work, one possible solution is to encode motion information in a different way so as to express subspace structure yet to preserve properties of articulated motion. Another approach consists in exploiting the inherent hierarchical structure of the articulated motion.

(a) Original frame corrupted by gaussian-distributed noise

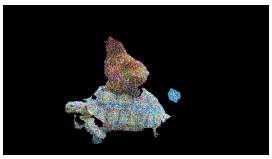

(b) Segmentation result

(c) Original frame corrupted by salt-and-pepper noise

(d) Segmentation result

Figure 6: Segmentation results on corrupted videos

5.2 Distributed optimization

Another possible extension of this work consists in taking advantage of the inherent parallel computing structure of ADMM. As demonstrated in the Algorithm 1, ADMM takes the form of a *decomposition-coordination* procedure, coordinating solutions to local subproblems to find a solution to a global problem. Therefore, it is natural to develop parallel and distributed optimization algorithm based on ADMM to solve efficiently the large-scale video co-segmentation problem.

A complete source code implementation of SSSC-AM in Python for reproducible research is available from the Authors at request.

References

[1] Pankaj K. Agarwal and Nabil H. Mustafa. k-means projective clustering. In *Proceedings of the Twenty-third ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '04, pages 155–165. ACM, 2004.

[2] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.

[3] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. *European Conference on Computer Vision*, 2010.

[4] J. Chang, Donglai Wei, and J.W. Fisher. A video representation using temporal superpixels. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2051–2058, June 2013.

[5] Wei-Chen Chiu and M. Fritz. Multi-class video co-segmentation with a generative multi-video model. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 321–328, June 2013.

[6] Ehsan Elhamifar and René Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2765–2781, 2013.

[7] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.

[8] Jiaming Guo, Loong-Fah Cheong, Robby T. Tan, and Steven Zhiying Zhou. Consistent foreground co-segmentation. *Computer Vision (ACCV)*, pages 241–257, 2014.

[9] Armand Joulin, Francis Bach, and Jean Ponce. Discriminative clustering for image co-segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1943–1950, 2010.

[10] Baohua Li, Ying Zhang, Zhouchen Lin, and Huchuan Lu. Subspace clustering by mixture of Gaussian regression. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[11] Chun-Guang Li and René Vidal. Structured sparse subspace clustering: A unified optimization framework. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[12] Chun-Guang Li, Chong You, and Rene Vidal. On sufficient condition for affine sparse subspace clustering. In *Signal Processing with Adaptive Sparse Structured Representations*, 2015.

[13] C. Liu, J. Yuen, and A. Torralba. Sift flow: Dense correspondence across scenes and its applications. In *IEEE Trans. Pattern Anal. Mach. Intell.*, 2011.

[14] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma. Robust recovery of subspace structures by low-rank representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):171–184, Jan 2013.

[15] Bryan Poling. A tutorial on camera models. University of Minnesota.

[16] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, Aug 2000.

[17] Mahdi Soltanolkotabi and Emmanuel J.Candés. A geometric analysis of subspace clustering with outliers. *The Annals of Statistics*, 40(4):2195–2238, 2012.

[18] Roberto Tron and René Vidal. A benchmark for the comparison of 3d motion segmentation algorithms. In *In CVPR*, 2007.

[19] R. Vidal, Yi Ma, and S. Sastry. Generalized principal component analysis (GPCA). *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(12):1945–1959, 2005.

[20] René Vidal. A tutorial on subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2010.

[21] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

[22] Chuan Wang, Yanwen Guo, Jie Zhu, Linbo Wang, and Wenping Wang. Video object co-segmentation via subspace clustering and quadratic pseudo-boolean optimization in an MRF framework. *Multimedia, IEEE Transactions on*, 16(4):903–916, June 2014.

[23] Wenguan Wang, Jianbing Shen, Xuelong Li, and Fatih Porikli. Robust video object cosegmentation. *IEEE Transactions on Image Processing*, 24(10):3137–3148, 2015.

[24] Dengsheng Zhang, Md. Monirul Islam, and Guojun Lu. A review on automatic image annotation techniques. *Pattern Recognition*, 45(1):346–362, 2012.

A Notation

A, B	Arbitrary $m \times n$ matrices
$a_{i,j}$	the (i,j)-th element of matrix A
a_i	The i-th column of A
$\text{vec}(A)$	$mn \times 1$ column vector obtained by stacking the columns of A on top of one another: $\text{vec}(A) = [a_{1,1}, \ldots, a_{m,1}, a_{1,2}, \ldots, a_{m,2}, \ldots, a_{1,n}, \ldots, a_{m,n}]^\top$
$\|A\|_1$	$\|A\|_1 = \sum_{i=1}^{m} \sum_{j=1}^{n} \|a_{ij}\|$
$\|A\|_{2,1}$	$\|A\|_{2,1} = \sum_{j=1}^{n} \|a_j\|_2 = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} \|a_{ij}\|^2\right)^{1/2}$
X	$n \times N$ data matrix. Each column refers to a data point.
Z	$N \times N$ coefficient matrix
$\text{diag}(Z)$	$\text{diag}(Z) \in \mathbb{R}^N$, vector of the diagonal elements of Z
E	$n \times N$ noise matrix
$\mathbf{0}$	Vector of appropriate dimensions filled with 0
$\mathbf{1}$	Vector of appropriate dimensions filled with 1
Q	$N \times k$ segmentation matrix with k the number of subspaces. $Q = [\mathbf{q_1}, \cdots, \mathbf{q_k}]$ indicates the membership of each data point to each subspace. $q_{i,j} = 1$ if the i-th column of X lies in subspace S_j and $q_{i,j} = 0$ otherwise.
$\boldsymbol{q}^{(i)}$	i-th row of matrix Q
$\mathcal{Q}$	Space of segmentation matrices: $\mathcal{Q} = \{Q \in \{0,1\}^{N \times k} : Q\mathbf{1} = \mathbf{1} \text{ and } \text{rank}(Q) = k\}$.
Θ	$N \times N$ matrix with $\Theta_{i,j} = \frac{1}{2} \left\|\boldsymbol{q}^{(i)} - \boldsymbol{q}^{(j)}\right\|^2$
$\odot$	Hadamard product or element-wise product: $(A \odot B)_{i,j} = (A)_{i,j} \cdot (B)_{i,j}$, $A, B \in \mathbb{R}^{m \times n}$
$\|Z\|_{1,Q}$	Subspace structured ℓ_1 norm: $\|Z\|_{1,Q} \doteq \|Z\|_1 + \alpha \|\Theta \odot Z\|_1 = \sum_{i,j} \|Z_{i,j}\|(1 + \frac{\alpha}{2} \left\|\boldsymbol{q}^{(i)} - \boldsymbol{q}^{(j)}\right\|^2)$, where $\alpha > 0$ is trade-off parameter.
$S_\tau(\nu)$	$S_\tau(\nu) = (\|\nu\| - \tau)_+ \text{sgn}(\nu)$
Affine subspace	A subset $U \subset V$ of a vector space V is an affine subspace if there exists a $u \in U$ such that $U - u = \{x - u \| x \in U\}$ is a linear subspace of V.
M	Motion matrix: $lM = \begin{bmatrix} x_{11} & \ldots & x_{1N} \\ \vdots & & \vdots \\ x_{F1} & \ldots & x_{FN} \end{bmatrix}_{2F \times N}$ where $\{x_{\text{fj}} \in \mathbb{R}^2\}_{j=1,\ldots,N}^{f=1,\ldots,F}$ denotes the 2D projection of N 3D points on rigidly moving objects onto F frames of a moving camera.